Managing Cisco NFVI Storage

This chapter describes basic architectural concepts that will help you understand the Cisco NFVI data storage architecture and data flow. It also provides techniques you can use to monitor the storage cluster health and the health of all systems that depend on it

Cisco NFVI Storage Architecture

OpenStack has multiple storage back ends. Cisco NFVI uses the Ceph back end. Ceph supports both block and object storage and is therefore used to store VM images and volumes that can be attached to VMs. Multiple OpenStack services that depend on the storage backend include:

• Glance (OpenStack image service)—Uses Ceph to store images.

• Cinder (OpenStack storage service)—Uses Ceph to create volumes that can be attached to VMs.

• Nova (OpenStack compute service)—Uses Ceph to connect to the volumes created by Cinder.

The following figure shows the Cisco NFVI storage architecture component model.

Figure 1: Cisco NFVI Storage Architecture
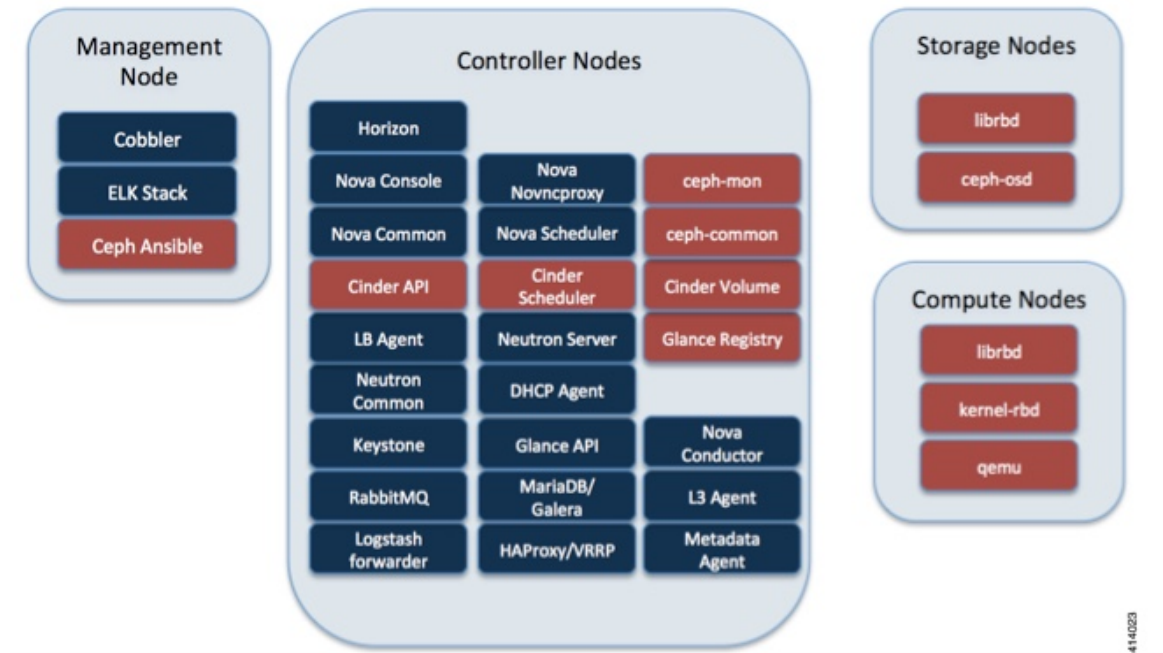
Verifying and Displaying Ceph Storage Pools

Ceph is configured with four independent pools: images, volumes, vms, and backups. (A default rbd pool is used internally.) Each Ceph pool is mapped to an OpenStack service. The Glance service stores data in the images pool, and the Cinder service stores data in the volumes pool. The Nova service can use the vms pool to boot ephemeral disks directly from the Ceph cluster depending on how the NOVA_BOOT_FROM option in the ~/openstack-configs/setup_data.yaml was configured prior to Cisco NFVI installation. If NOVA_BOOT_FROM is set to ceph before you run the Cisco NFVI installation, the Nova service will boot from the Ceph vms pool. By default, NOVA_BOOT_FROM is set to local, which means that all VM ephemeral disks are stored as files in the compute nodes. Changing this option after installation does not affect the use of the vms pool for ephemeral disks.

The Glance, Cinder, and Nova OpenStack services depend on the Ceph cluster for backend storage. Therefore, they need IP connectivity to the controller nodes. The default port used to connect Glance, Cinder, and Nova to the Ceph cluster is 6789. Authentication through cephx is required, which means authentication tokens, called keyrings, must be deployed to the OpenStack components for authentication.

To verify and display the Cisco NFVI Ceph storage pools:

Step 1 Launch a SSH session to a controller node, for example:
```
[root@management-server-cisco ~]# ssh root@controller_server-1
```
Step 2 Navigate to the Ceph Monitor container:
```
[root@controller_server-1 ~]# cephmon
```
Step 3 List the Ceph pools:
```
cephmon_4612 [root@controller_server-1 ~]# ceph osd lspools
0 rbd,1 images,2 volumes,3 vms,4 backups,
```
Step 4 List the images pool content:
```
cephmon_4612 [ceph@controller_server-1 /]$ rbd list images
a4963d51-d3b7-4b17-bf1e-2ebac07e1593
```

Checking the Storage Cluster Health

Cisco recommends that you perform a few verifications to determine whether the Ceph cluster is healthy and is connected to the Glance, Cinder, and Nova OpenStack services, which have Ceph cluster dependencies. The first task to check the health of the cluster itself by completing the following steps:

Step 1 From the Cisco NFVI management node, launch a SSH session to a controller node, for example:
```
[root@management-server-cisco ~]# ssh root@controller_server-1
```
Step 2 Navigate to the Ceph Monitor container:
```
[root@controller_server-1 ~]# cephmon
```
Step 3 Check the Ceph cluster status:
```
cephmon_4612 [ceph@controller_server-1 ceph]$ ceph status
```
Sample response:
```
cluster dbc29438-d3e0-4e0c-852b-170aaf4bd935
    health HEALTH_OK
    monmap e1: 3 mons at {ceph-controller_server-1=20.0.0.7:6789/0,
  ceph-controller_server-2=20.0.0.6:6789/0,ceph-controller_server-3=20.0.0.5:6789/0}
          election epoch 8, quorum 0,1,2 ceph-controller_server-3,
  ceph-controller_server-2,ceph-controller_server-1
    osdmap e252: 25 osds: 25 up, 25 in
     pgmap v593: 1024 pgs, 5 pools, 406 MB data, 57 objects
          2341 MB used, 61525 GB / 61527 GB avail
              1024 active+clean
```
This example displays three monitors, all in good health, and 25 object storage devices (OSDs). All OSDs show as up and in the cluster.

Step 4 To see a full listing of all OSDs sorted by storage node, enter:
```
cephmon_4612 [ceph@controller_server-1 ceph]$ ceph osd tree
```

Sample response:

```
ID WEIGHT    TYPE NAME               UP/DOWN REWEIGHT PRIMARY-AFFINITY
-1 60.18979 root default
-2 18.96994     host controller_server-2
 1  2.70999         osd.1              up  1.00000          1.00000
 5  2.70999         osd.5              up  1.00000          1.00000
 6  2.70999         osd.6              up  1.00000          1.00000
11  2.70999         osd.11             up  1.00000          1.00000
12  2.70999         osd.12             up  1.00000          1.00000
17  2.70999         osd.17             up  1.00000          1.00000
20  2.70999         osd.20             up  1.00000          1.00000
-3 18.96994     host controller_server-1
 0  2.70999         osd.0              up  1.00000          1.00000
 4  2.70999         osd.4              up  1.00000          1.00000
 8  2.70999         osd.8              up  1.00000          1.00000
10  2.70999         osd.10             up  1.00000          1.00000
13  2.70999         osd.13             up  1.00000          1.00000
16  2.70999         osd.16             up  1.00000          1.00000
18  2.70999         osd.18             up  1.00000          1.00000
-4 18.96994     host controller_server-3
 2  2.70999         osd.2              up  1.00000          1.00000
 3  2.70999         osd.3              up  1.00000          1.00000
 7  2.70999         osd.7              up  1.00000          1.00000
 9  2.70999         osd.9              up  1.00000          1.00000
14  2.70999         osd.14             up  1.00000          1.00000
15  2.70999         osd.15             up  1.00000          1.00000
19  2.70999         osd.19             up  1.00000          1.00000
-5  3.27997     host controller_server-4
21  0.81999         osd.21             up  1.00000          1.00000
22  0.81999         osd.22             up  1.00000          1.00000
23  0.81999         osd.23             up  1.00000          1.00000
24  0.81999         osd.24             up  1.00000          1.00000
```

What to Do Next

After you verify the Ceph cluster is in good health, check that the individual OpenStack components have connectivity and their authentication tokens—keyrings—match the Ceph Monitor keyrings. The following procedures show how to check the connectivity and authentication between Ceph and Glance, Ceph and Cinder, and Ceph and Nova.

Checking Glance Connectivity

The Glance API container must be connected to the Cisco NFVI controller nodes. Complete the following steps to verify the Glance to controller node connectivity:

Step 1   From the management node, launch a SSH session to a controller node, for example:

```
[root@management-server-cisco ~]# ssh root@controller_server-1
```

Step 2   Navigate to the Glance API container:

```
[root@controller_server-1 ~]# glanceapi
```

Step 3   Check the Glance API container connectivity to a controller node different from the one entered in Step 1, in this case, controller_server 2:

```
glanceapi_4612 [glance@controller_server-1 /]$ curl controller_server-2:6789
```
If the connection is successful, you will see a message like the following:

```
glanceapi_4612 [glance@controller_server-1 /]$ curl controller_server-2:6789
ceph v027?
```
If the connection is not successful, you will see a message like the following:

```
glanceapi_4612 [glance@controller_server-1 /]$ curl controller_server-2:6789
curl: (7) Failed connect to controller_server-2:6789; Connection refused
```
A message like the one above means the Ceph monitor running on the target controller node controller_server-2 is not listening on the specified port or there is no route to it from the Glance API container.

Checking one controller node should be enough to ensure one connection path available for the Glance API. However, because Cisco NFVI controller nodes run as part of an HA cluster, you should run Step 3 above targeting all the controller nodes in the Cisco NFVI pod.

What to Do Next

After you verify the Glance API connectivity to all Cisco NFVI controller nodes, check the Glance keyring to ensure it matches the Ceph monitor keyring.

Verifying Glance and Ceph Monitor Keyrings

Complete the following steps to verify the Glance API keyring matches the Ceph Monitor keyring.

Step 1   Launch a SSH session to a controller node, for example:

```
[root@management-server-cisco ~]# ssh root@controller_server-1
```

Step 2   Navigate to the Glance API container:

```
[root@controller_server-1 ~]# glanceapi
```

Step 3   Check the Glance keyring content, for example:

```
glanceapi_4612 [glance@controller_server-1 /]$ cat /etc/ceph/client.glance.keyring
[client.glance]
key = AQA/pY1XBAnHMBAAeS+0Wmh9PLZe1XqkIW/p0A==
```

Step 4   Navigate to the Ceph Monitor container:

```
[root@controller_server-1 ~]# cephmon
```

Step 5   Display the Ceph Monitor keyring content:

```
cephmon_4612 [ceph@controller_server-1 ceph]$ cat /etc/ceph/ceph.client.glance.keyring
[client.glance]

key = AQA/pY1XBAnHMBAAeS+0Wmh9PLZe1XqkIW/p0A==
```
Verify the keyring matches the Glance API keyring displayed in Step 3.

What to Do Next

A final check to ensure that Ceph and Glance are connected is to actually import a Glance image using Horizon or the Glance CLI. After you import an image, compare the IDs seen by Glance and by Ceph. They should match, indicating Ceph is handling the backend for Glance.

Verifying Glance Image ID on Ceph

The following steps verify Ceph is properly handling new Glance images by checking that the image ID for a new Glance image is the same as the image ID displayed in Ceph.

Step 1   From the management node, load the OpenStack authentication variables:
```
[root@management-server-cisco ~]# source ~/openstack-configs/openrc
```

Step 2   Import any Glance image. In the example below, a RHEL 7.1 qcow2 image is used.
```
[root@management-server-cisco images]# glance image-create
--name "rhel" --disk-format qcow2  --container-format bare --file
rhel-guest-image-7.1-20150224.0.x86_64.qcow2
```

Step 3   List the Glance images:
```
[root@management-server-cisco images]# glance image-list | grep rhel
| a4963d51-d3b7-4b17-bf1e-2ebac07e1593 | rhel
```

Step 4   Navigate to the Ceph Monitor container:
```
[root@controller_server-1 ~]# cephmon
```

Step 5   Display the contents of the Ceph images pool:
```
cephmon_4612 [ceph@controller_server-1 ceph]$ rbd list images | grep
a4963d51-d3b7-4b17-bf1e-2ebac07e1593
a4963d51-d3b7-4b17-bf1e-2ebac07e1593
```

Step 6   Verify that the Glance image ID displayed in Step 3 matches the image ID displayed by Ceph.

Checking Cinder Connectivity

The Cinder volume container must have connectivity to the Cisco NFVI controller nodes. Complete the following steps to verify Cinder volume has connectivity to the controller nodes:

Step 1   From the management node, launch a SSH session to a controller node, for example:
```
[root@management-server-cisco ~]# ssh root@controller_server-1
```

Step 2   Navigate to the Cinder volume container:
```
[root@controller_server-1 ~]# cindervolume
```

Step 3   Check the Cinder volume container connectivity to a controller node different from the one entered in Step 1, in this case, controller_server-2:

```
cindervolume_4612 [cinder@controller_server-1 /]$ curl controller_server-2:6789
```

If the connection is successful, you will see a message like the following:

```
cindervolume_4612 [cinder@controller_server-1 /]$ curl controller_server-2:6789
ceph v027?
```

If the connection is not successful, you will see a message like the following:

```
cindervolume_4612 [cinder@controller_server-1 /]$ curl controller_server-2:6789
curl: (7) Failed connect to controller_server-2:6789; Connection refused
```

A message like the one above means the Ceph monitor running on the target controller node controller_server-2 is not listening on the specified port or there is no route to it from the Cinder volume container.

Checking one controller node should be enough to ensure one connection path is available for the Cinder volume. However, because Cisco NFVI controller nodes run as part of an HA cluster, repeat Step 3 targeting all the controller nodes in the Cisco NFVI pod.

What to Do Next

After you verify the Cinder volume connectivity to all Cisco NFVI controller nodes, check the Cinder keyring to ensure it matches the Ceph monitor keyring.

Verifying the Cinder and Ceph Monitor Keyrings

Complete the following steps to verify the Cinder volume keyring matches the Ceph Monitor keyring.

Step 1   From the management node, launch a SSH session to a controller node, for example:

```
[root@management-server-cisco ~]# ssh root@controller_server-1
```

Step 2   Navigate to the Cinder volume container:

```
[root@controller_server-1 ~]# cindervolume
```

Step 3   Check the Cinder keyring content, for example:

```
cindervolume_4612 [cinder@controller_server-1 /]$ cat /etc/ceph/client.cinder.keyring
[client.cinder]
key = AQA/pY1XBAnHMBAAeS+0Wmh9PLZe1XqkIW/p0A==
```

Step 4   Navigate to the Ceph Monitor container:

```
[root@controller_server-1 ~]# cephmon
```

Step 5   Display the Ceph Monitor keyring content:

```
cephmon_4612 [ceph@controller_server-1 ceph]$ cat /etc/ceph/ceph.client.cinder.keyring
[client.cinder]

key = AQA/pY1XBAnHMBAAeS+0Wmh9PLZe1XqkIW/p0A==
```

Verify the keyring matches the Cinder volume keyring displayed in Step 3.

What to Do Next

As a final Ceph and Cinder connectivity verification, import a Cinder image using Horizon or the Cinder CLI. After you import the image, compare the IDs seen by Cinder and by Ceph. They should match, indicating Ceph is handling the backend for Cinder.

Verifying the Cinder Volume ID on Ceph

The following steps verify Ceph is properly handling new Cinder volumes by checking that the volume ID for a new Cinder volume is the same as the volume ID displayed in Ceph.

Step 1 From the management node, load the OpenStack authentication variables:

```
[root@management-server-cisco ~]# source ~/openstack-configs/openrc
```

Step 2 Create an empty volume:

```
[root@management-server-cisco ~]# cinder create --name ciscovol1 5
```
The above command will create a new 5 GB Cinder volume named ciscovol1.

Step 3 List the Cinder volumes:

```
[[root@management-server-cisco ~]# cinder list
+------------------------------------+-------+-----------------+--
|                 ID                 | Status | Migration Status |...

| dd188a5d-f822-4769-8a57-c16694841a23 | in-use |        -         |...
+----------------------------------+------+
```

Step 4 Navigate to the Ceph Monitor container:

```
[root@controller_server-1 ~]# cephmon
```

Step 5 Display the contents of the Ceph volumes pool:

```
cephmon_4612 [ceph@controller_server-1 ceph]$ rbd list volumes
volume-dd188a5d-f822-4769-8a57-c16694841a23
```

Step 6 Verify that the Cinder volume ID displayed in Step 3 matches the volume ID displayed by Ceph, excluding the "volume-" prefix.

Checking Nova Connectivity

The Nova libvirt container must have connectivity to the Cisco NFVI controller nodes. Complete the following steps to verify Nova has connectivity to the controller nodes:

Step 1 From the management node, launch a SSH session to a controller node, for example:

```
[root@management-server-cisco ~]# ssh root@Computenode_server-1
```

Step 2 Navigate to the Nova libvirt container:

```
[root@compute_server-1 ~]# libvirt
```

Step 3 Check the Nova libvirt container connectivity to a controller node, in this case, controller_server 1:
```
novalibvirt_4612 [root@compute_server-1 /]$ curl controller_server-2:6789
```
If the connection is successful, you will see a message like the following:
```
novalibvirt_4612 [root@compute_server-1 /]$ curl controller_server-1:6789
ceph v027?
```
If the connection is not successful, you will see a message like the following:
```
novalibvirt_4612 [root@compute_server-1 /]$ curl controller_server-1:6789
curl: (7) Failed connect to controller_server-1:6789; Connection refused
```
A message like the one above means the Ceph monitor running on the target controller node controller_server-1 is not listening on the specified port or there is no route to it from the Nova libvirt container.

Checking one controller node should be enough to ensure one connection path available for the Nova libvirt. However, because Cisco NFVI controller nodes run as part of an HA cluster, you should run Step 3 above targeting all the controller nodes in the Cisco NFVI pod.

What to Do Next

After you verify the Nova libvirt connectivity to all Cisco NFVI controller nodes, check the Nova keyring to ensure it matches the Ceph monitor keyring.

Verifying the Nova and Ceph Monitor Keyrings

Complete the following steps to verify the Nova libvert keyring matches the Ceph Monitor keyring.

Step 1 From the management node, launch a SSH session to a controller node, for example:
```
[root@management-server-cisco ~]# ssh root@controller_server-1
```
Step 2 Navigate to the Nova libvert container:
```
[root@compute_server-1 ~]# libvirt
```
Step 3 Extract the libvirt secret that contains the Nova libvirt keyring:
```
novalibvirt_4612 [root@compute_server-1 /]# virsh secret-list
UUID                                      Usage ...
----------------------------------------------------------------
b5769938-e09f-47cb-bdb6-25b15b557e84   ceph client.cinder ...
```
Step 4 Get the keyring from the libvert secret:
```
novalibvirt_4612 [root@controller_server-1 /]# virsh secret-get-value
b5769938-e09f-47cb-bdb6-25b15b557e84
AQBApY1XQCBBEBAAroXvmiwmlSMEyEoXKl/sQA==
```
Step 5 Navigate to the Ceph Monitor container:
```
[root@controller_server-1 ~]# cephmon
```
Step 6 Display the Ceph Monitor keyring content:
```
cephmon_4612 [ceph@controller_server-1 ceph]$ cat /etc/ceph/ceph.client.cinder.keyring
[client.cinder]

key = AQBApY1XQCBBEBAAroXvmiwmlSMEyEoXKl/sQA==
```

Verify the keyring matches the Nova libvert keyring displayed in Step 3. Notice that in the above example the Cinder keyring is checked even though this procedure is for the Nova libvert keyring. This occurs because the Nova services need access to the Cinder volumes and so authentication to Ceph uses the Cinder keyring.

What to Do Next

Complete a final check to ensure that Ceph and Nova are connected by attaching a Nova volume using Horizon or the Nova CLI. After you attach the Nova volume, check the libvert domain.

Verifying Nova Instance ID

From the management node, complete the following steps to verify the Nova instance ID:

Step 1 Load the OpenStack authentication variables:

```
[root@management-server-cisco installer]# source ~/openstack-configs/openrc
```

Step 2 List the Nova instances:

```
[root@management-server-cisco images]# nova list
+--------------------------------------+-----------+--------+--------
| ID                                   | Name      | Status | Task
+--------------------------------------+-----------+--------+--------
| 77ea3918-793b-4fa7-9961-10fbdc15c6e5 | cisco-vm  | ACTIVE | -
+--------------------------------+-----------+-------+-
```

Step 3 Show the Nova instance ID for one of the instances:

```
[root@management-server-cisco images]# nova show
77ea3918-793b-4fa7-9961-10fbdc15c6e5 | grep instance_name
| OS-EXT-SRV-ATTR:instance_name      | instance-00000003
```

The Nova instance ID in this example is instance-00000003. This ID will be used later with the virsh command. Nova instance IDs are actually the libvirt IDs of the libvirt domain associated with the Nova instance.

Step 4 Identify the compute node where the VM was deployed:

```
[root@management-server-cisco images]# nova show 77ea3918-793b-4fa7-9961-10fbdc15c6e5 | grep hypervisor
| OS-EXT-SRV-ATTR:hypervisor_hostname | compute_server-1
```

The compute node in this case is compute_server-1. You will connect to this compute node to call the virsh commands. Next, you get the volume ID from the libvirt domain in the Nova libvirt container.

Step 5 Launch a SSH session to the identified compute node, compute_server-1:

```
[root@management-server-cisco ~]# ssh root@compute_server-1
```

Step 6 Navigate to the Nova libvirt container:

```
[root@compute_server-1 ~]# libvirt
```

Step 7 Get the instance libvirt domain volume ID:

```
novalibvirt_4612 [root@compute_server-1 /]# virsh dumpxml instance-00000003 | grep rbd
 <source protocol='rbd' name='volumes/volume-dd188a5d-f822-4769-8a57-c16694841a23'>
```

Step 8 Launch a SSH session to a controller node:

```
[root@management-server-cisco ~]# ssh root@controller_server-1
```

Step 9   Navigate to the Ceph Monitor container:

`[root@compute_server-1 ~]# cephmon`

Step 10   Verify volume ID matches the ID in Step 7:

```
cephmon_4612 [ceph@controller_server-1 ceph]
$ rbd list volumes | grep volume-dd188a5d-f822-4769-8a57-c16694841a23
volume-dd188a5d-f822-4769-8a57-c16694841a23
```

Displaying Docker Disk Space Usage

Docker supports multiple storage back ends such as Device Mapper, thin pool, overlay, and AUFS. Cisco VIM uses the devicemapper storage driver because it provides strong performance and thin provisioning. Device Mapper is a kernel-based framework that supports advanced volume management capability. Complete the following steps to display the disk space used by Docker containers.

Step 1   Launch a SSH session to a controller or compute node, for example:

`[root@management-server-cisco ~]# ssh root@controller_server-1`

Step 2   Enter the docker info command to display the disk space used by Docker containers:

```
[root@controller_server_1 ~]# docker info
Containers: 24
Images: 186
Storage Driver: devicemapper
 Pool Name: vg_var-docker--pool
 Pool Blocksize: 524.3 kB
 Backing Filesystem: xfs
 Data file:
 Metadata file:
 Data Space Used: 17.51 GB
 Data Space Total: 274.9 GB
 Data Space Available: 257.4 GB...
```

Reconfiguring SwiftStack Integration

Cisco VIM provides integration with SwiftStack, an object storage solution. The key aspect of the SwiftStack integration is to add a SwfitStack endpoint to an existing pod running on Cisco VIM through the reconfigure option. In this case the SwiftStack is installed and managed outside the Cisco VIM ahead of time, and the VIM orchestrator adds the relevant Keystone configuration details to access the SwiftStack endpoint (see the Cisco VIM install guide for more details of SwiftStack).

The following options support the SwiftStack reconfiguration:

  • Enable SwiftStack integration if it is not present.

  • Reconfigure the existing SwiftStack PAC endpoint to point to a different cluster (cluster_api_endpoint).

• Reconfigure the Reseller_prefix of the existing SwiftStack installation.

• Reconfigure the admin password (admin_password) of an existing SwiftStack Install.

Integrating SwiftStack over TLS

The automation supports SwiftStack integration over TLS. To enable TLS, the CA root certificate must be presented as part of the `/root/openstack-configs/haproxy-ca.crt` file. The protocol parameter within the SWIFTSTACK stanza must be set to https. As a pre-requisite, the SwiftStack cluster needs to be configured to enable HTTPS connections for the SwiftStack APIs with termination at the proxy servers.

The following section needs to be configured in the Setup_data.yaml file.

```
###########################################
# Optional Swift configuration section
###########################################
# SWIFTSTACK: # Identifies the objectstore provider by name
#    cluster_api_endpoint: <IP address of PAC (proxy-account-container) endpoint>
#    reseller_prefix: <Reseller_prefix as configured for Keysone Auth,AuthToken support in
Swiftstack E.g KEY_>
#    admin_user: <admin user for swift to authenticate in keystone>
#    admin_password: <swiftstack_admin_password>
#    admin_tenant: <The service tenant corresponding to the Account-Container used by
Swiftstack
     protocol: <http or https> # protocol that swiftstack is running on top
```

Note The operator should pay attention while updating the settings to ensure that SwiftStack over TLS are appropriately pre-configured in the customer-managed SwiftStack controller as specified in the Install guide.

To initiate the integration, copy the `setupdata` into a local directory by running the following command:

```
[root@mgmt1 ~]# cd /root/
[root@mgmt1 ~]# mkdir MyDir
[root@mgmt1 ~]# cd MyDir
[root@mgmt1 ~]# cp /root/openstack-configs/setup_data.yaml <my_setup_data.yaml>
```

Update the `setupdata` by running the following command:

```
[root@mgmt1 ~]# vi my_setup_data.yaml (update the setup_data to include SwiftStack info)
```

Run the reconfiguration command as follows:

```
[root@mgmt1 ~]# cd ~/installer-xxxx
[root@mgmt1 ~]# ciscovim --setupfile ~/MyDir/<my_setup_data.yaml> reconfigure
```

Cinder Volume Backup on SwiftStack

Cisco VIM enables cinder service to be configured to backup its block storage volumes to the SwiftStack object store. This feature is automatically configured if the SWIFTSTACK stanza is present in the setup_data.yaml file. The mechanism is to authenticate against SwiftStack during volume backups leverages. The same keystone SwiftStack endpoint is configured to manage objects. The default SwiftStack container that manages cinder volumes within the account (Keystone Tenant as specified by admin_tenant) is currently defaulted to volumebackups.

Reconfiguring Administrator Source Networks

To access the administrator services, Cisco VIM 2.2 provides source IP based filtering of network requests on the management node. These services include SSH and Kibana dashboard access. When the services are configured all admin network requests made to the management node are dropped, except those from white listed addresses in the configuration.

Reconfiguring administrator source network supports the following options:

- Set administrator source network list: Network addresses can be added or deleted from the configuration; the list is replaced in whole during a reconfigure operation.

- Remove administrator source network list: If the **admin_source_networks** option is removed, then the source address will not filter the incoming admin service requests.

The following section needs to be configured in the Setup_data.yaml file:

```
admin_source_networks: # optional, host based firewall to white list admin's source IP
  - 10.0.0.0/8
  - 172.16.0.0/12
```

Note
The operator should to be careful while updating the source networks. If the list is mis-configured, operators may lock themselves out of access to the management node through SSH. If this happens, an operator must log into the management node through the console port to repair the configuration.

To initiate the integration, copy the `setupdata` into a local directory by running the following command:

```
[root@mgmt1 ~]# cd /root/
[root@mgmt1 ~]# mkdir MyDir
[root@mgmt1 ~]# cd MyDir
[root@mgmt1 ~]# cp /root/openstack-configs/setup_data.yaml <my_setup_data.yaml>
```

Update the `setupdata` by running the following command:

```
[root@mgmt1 ~]# vi my_setup_data.yaml (update the setup_data to include SwiftStack info)
```

Run the reconfiguration command as follows:

```
[root@mgmt1 ~]# cd ~/installer-xxxx
[root@mgmt1 ~]# ciscovim --setupfile ~/MyDir/<my_setup_data.yaml> reconfigure
```

Password Reset for Cisco VIM Management Node

Run the following command to reset the Root Password of Cisco VIM management node **RHEL-7 / systemd** .

1 Boot your system and wait until the **GRUB2** menu appears.

2 In the **boot loader** menu, highlight any entry and press **e**.

3 Find the line beginning with linux. At the end of this line, append the following:

```
 init=/bin/sh
```
Or if you face any alarm, instead of **ro** change **rw** to **sysroot** as shown in the following example:

```
rw init=/sysroot/bin/sh
```
4 Press **Ctrl+X** to boot the system using the options you just edited.

Once the system boots, you will be presented with a shell prompt without having to enter any user name or password:

```
sh-4.2#
```

5 Load the installed SELinux policy by running the following command:

```
sh-4.2# /usr/sbin/load_policy -i
```

6 Execute the following command to remount your root partition:

```
sh4.2#
    mount -o remount,rw /
```

7 Reset the root password by running the following command:

```
sh4.2# passwd root
```
When prompted, enter your new root password and confirm by pressing the **Enter** key. Enter the password for the second time to make sure you typed it correctly and confirm with **Enter** again. If both the passwords match, a message informing you of a successful root password change will appear.

8 Execute the following command to remount the root partition again, this time as read-only:

```
sh4.2#
    mount -o remount,ro /
```

9 Reboot the system. Now you will be able to log in as the root user using the new password set up during this procedure.

To reboot the system, enter exit and exit again to leave the environment and reboot the system.

References: https://access.redhat.com/solutions/918283.